Analysis of alerting system failures in commercial aviation accidents

Randall J. Mumaw
NASA Ames Research Center

The role of an alerting system is to make the system operator (e.g., pilot) aware of an impending hazard or unsafe state so the hazard can be avoided or managed successfully. A review of 46 commercial aviation accidents (between 1998 and 2014) revealed that, in the vast majority of events, either the hazard was not alerted or relevant hazard alerting occurred but failed to aid the flight crew sufficiently. For this set of events, alerting system failures were placed in one of five phases: Detection, Understanding, Action Selection, Prioritization, and Execution. This study also reviewed the evolution of alerting system schemes in commercial aviation, which revealed naïve assumptions about pilot reliability in monitoring flight path parameters; specifically, pilot monitoring was assumed to be more effective than it actually is. Examples are provided of the types of alerting system failures that have occurred, and recommendations are provided for alerting system improvements.

The Role of Alerting

In any complex, safety-critical system, alerting should initiate the appropriate operator response to managing a hazard (or, more generally, an unsafe state). Specifically, the alert and the other elements of the system interface need to help the operator

- Orient to an important change – Some salient stimulus, such as a loud sound/voice or bright light, or both, needs to attract the operator's attention that some important change has occurred requiring the operator's immediate awareness.
- Understand the nature and urgency of that change – Some interface element needs to describe this important change. Information about the nature of the change, as well as the urgency and/or the importance of response need to be presented clearly and simply.
- Identify appropriate actions – When there are operator actions that are recommended to manage the failure or hazard, those actions have to be presented to the operator (although in some cases, an action or small set of actions is memorized by the operator). The alert needs to create a link to those actions, either by making them part of the alert or by pointing to a set of actions.
- Identify the priority for the actions – In some non-normal situations, there may be a number of actions for the operator to take, especially when there are multiple alerts. The system interface needs to aid the operator in assessing which actions have the highest priority (or the highest urgency).
- Execute the actions efficiently, accurately, and completely – The system interface and the operational documents need to support the operator in executing the prescribed actions. In addition, the interface should provide clear feedback that the unsafe state has been resolved (or cannot be resolved).

These five phases represent an **integrated alerting sequence**, which can be used to identify the point at which alerting fails when it does not lead to successfully managing the situation. For example, a flight crew may fail to detect an alert (perhaps because it was insufficiently salient), which would preclude any further steps to manage the hazard.

How Effective is Alerting?

The design of an alerting system is evaluated during system (airplane) certification. From my experience, that evaluation is strongly driven by subjective reviews from pilots/SMEs (vs empirical studies). Further, these reviews place significant value on the mere presence of a hazard alert. There is little formal exploration of the larger alerting sequence.

Operational experience, although slow to accrue, is another measure of effectiveness, and reviews of safety events can reveal more about the performance of alerting systems. In 2010, the Commercial Aviation Safety Team (CAST) formed a team to review a set of 18 accidents and incidents tied to Loss of Control (LOC). This review revealed that, in every event, alerting had failed to support the flight crew in successfully managing the hazard (CAST, 2014). The question this raised was, "Was this relatively small sample representative of the full range of accidents?"

The goal in the present study was to review a large, representative set of commercial aviation accidents to determine whether the alerting system was effective, and if it was not, in what way it failed. [Note that accident causality is typically complex and there is rarely (almost never) a single cause. The discussion here regards the role that the alerting system did play as a contributor to the event vs what role it could have played.]

Method: Event Set

A review was conducted of commercial aviation accidents (defined here as events with on-board fatalities) and seven major incidents (no fatalities but significant upset) to determine the effectiveness of the alerting system. Initially, a list of all accidents in large, western-built, commercial transports from 2002 to 2014 was developed (plus three accidents from between 1998 and 2002 that had been analyzed by the CAST effort). Events were drawn from the event database at https://aviation-safety.net. Events were excluded from this set for one of three reasons: terrorism, pilot suicide, or the accident report was incomplete or unavailable.

There were 132 accidents in the resulting set, and with the seven incidents, the total event set was 139 (representing 6,052 on-board fatalities). The largest commercial transports—Boeing and Airbus airplanes and the more recent McDonnell-Douglas airplanes (DC-9s, MD-80s, MD-11s)—made up slightly more than half of these events, accounting for 74 of the 139 (53%) events.

Out of this set of 139, 64 (57 accidents and 7 incidents) were selected for review and analysis. Preference was given to the largest airplanes, which create the largest exposure for the flying public. This set of 64 represents 46% of the complete set of events. If one considers just the largest transports (Boeing, Airbus, and more-recent McDonnell-Douglas products), 57 of those 74 events (or 77%) were selected for analysis. There were two additional reasons for sampling primarily from these fleets:

1. These airplanes are more likely to have more complete and better-integrated alerting systems,
2. The accident reports are more likely to be thorough and more easily accessed.

The 64 selected events represent 4,020 fatalities, which is 66% of the total fatalities in the 139 events.

Method: Analysis Procedure

For this study, a hazard is defined as an undesirable airplane condition/state that has been alerted in some airplane designs. The hazards considered for this study are the following

- Bank angle – The airplane rolled (or was rolled) beyond 45° into a loss of control (LOC) situation that was not recovered.
- Low airspeed – Airspeed dropped below normal operating speeds.
- Approach to stall – Airspeed dropped even lower and/or airplane angle of attack increased to a point where the airplane was close to stalling (or stalls).
- Ground proximity – The airplane flew near terrain (when it was not landing).
- Landing configuration / High-energy approach – The airplane landed fast and/or the approach was above the glidepath or fast, and it became very difficult to manage the landing.
- Take-off configuration – The airplane was not properly configured for take-off, leading to insufficient lift to fly at the expected take-off speed.
- Cabin altitude – Pressure was not maintained in the airplane as it climbed, and the airplane could not sustain human life.
- Unreliable airspeed – Airspeed indication became erroneous due to a fault in air data; at least one of the airspeed indicators became invalid.
- Take-off performance – There was insufficient thrust for take-off with the available runway length.
- Autothrottle disengagement – The autothrottle disengaged (without pilot input) and the pilot failed to manage thrust.
- Engine out – Loss of an engine.

For each of the 64 events, I identified the hazards that were present (from the list above), the hazards that led to the accident, the airplane alerts that occurred, and the flight crew's response to any alerting (a description of their actions compared to the actions that would have been expected). These data are taken from the descriptions in the accident report. In some accidents, multiple hazards—e.g., low airspeed and approach to stall—were present in the event. There was no re-analysis of the accident; that is, the information in the accident report was treated as accurate.

More specifically, the analysis determined if the alerting system failed in one of the following ways:

- Alert Absent / Fail to detect the alert or hazard – Many of these were cases in which there was no alert for the hazard (or the alert malfunctioned). Or, the flight crew or the pilot was seemingly unaware that an alert had occurred (their behavior did not reveal an awareness). Note that Table 1 splits out these into two categories.
- Alert not understood – These were cases in which the flight crew or pilot was aware there had been an

Hazard	No Alert	Detect	Understand	Select Action	Prioritize	Execution	
Low Airspeed	8	3					11
Approach to Stall				9		4	13
Bank Angle	1		1	4			6
Unrel Airspeed			1	2			3
Ground Prox	4	5	3				12
Cabin Altitude			1				1
High-energy Approach	1	2					3
A/T Disengage		1					1
TO Performance	2						2
TO Configuration	1						1
Engine Out				1			1
	17	11	6	16	0	4	54

Table 1. Categorization of Alerting Failures (Note: some events had two hazards to alert)

alert, but they did not understand what the hazard was.
- Fail to select the appropriate action – These were cases in which the flight crew or pilot took inappropriate actions—actions different from what would be expected in the hazard situation. In some cases, these were opposite actions, such as roll the wrong way. In other cases, the flight crew had not been trained on the correct action.
- Fail to assign high priority to the alert – These were cases in which the alert or the alert response was not given sufficiently high priority.
- Fail to execute the appropriate response – These were cases where the flight crew or pilot understood what the response should be but their response was poorly performed.

The categorization of each event is described in the more complete technical report (Mumaw, 2017).

Results

[Note that the analysis, at the time of this writing (June, 2017), was complete for 46 of the 64 accidents; the complete analysis appears in a NASA technical report.]

Table 1 shows the resulting categorization of alerting system failures in 42 events [Note that, of the 46 events analyzed, only four had no alerting issues!]. The following paragraphs break out the types of failures encountered for some of the hazards (when there were multiple events).

There were 11 low-airspeed situations in this set. For seven of the events (five accidents and two incidents), there was no alerting for the low-airspeed situation. In another event, the alert malfunctioned and did not occur. And, in the last three events, the visual-only alert was not detected by the flight crew. For all of these, alerting failed because there was no alert, or the alert was not detected.

Nine of these 11 events transitioned from low airspeed to an approach-to-stall situation and alert (and there were four other events that had a stick shaker for a total of 13). These failed in a different way. For nine of these 13, the pilot flying (PF) responded to the stick shaker in a manner quite different from the way he should have responded. In some of these cases, this may have been tied to inadequate training on these types of situations. In two cases, there was no training at all on responding to a stall warning. In two of the remaining four cases, the pilot inadvertently got into an autoflight mode that prevented an effective response. In the last two cases, the pilot was in a difficult-to-manage situation and handled it poorly. So, unlike low airspeed, of which the pilot was unaware, stick shaker made the pilot aware of the situation. He just did not handle that situation well enough to recover.

Bank angle accidents were more diverse; there were six events. In one accident, there was no alert. In another accident, an alert came on, was silenced by one of the pilots, and then was unheeded. Perhaps the pilots did not understand the alert. In the last four events, the pilot did not understand what actions fit the situation and rolled the airplane away from wings-level and into a loss of control.

There were three events involving unreliable airspeed, two incidents and one accident. In one of these, the flight crew did not understand the situation and chased the bad airspeed. In the other two events, the

crew knew that they had airspeed issues but failed to link to the appropriate checklist.

The ground proximity hazard led to 12 events. As with low airspeed, some of these went unalerted. Older Ground Proximity Warning Systems (GPWS) failed to pick up approaching terrain in certain situations and failed to alert. A more interesting finding are the five events in which the flight crew kept flying toward the terrain even though the GPWS alert was active. This is perhaps due to a cognitive tunneling (or fixation) that seemed to be a failure to detect the alert (Dehais et al., 2014). Finally, in three other cases, the crew seemed aware of the GPWS alert but made an inadequate response to it; perhaps not fully understanding where they were relative to the terrain.

The high-energy approach hazard showed some similar failures. For two of these three events, the flight crew continued on despite continuous alerting (cognitive tunneling, or perhaps a conscious decision to ignore the alert because they thought they were managing it). For the other event, it got into a landing configuration condition that was not alerted.

Discussion

The analysis presented here demonstrates that the presence of a hazard alert is not always sufficient to initiate an effective response. To be fair, this analysis focused on only one leg of the three-legged stool: alerting system design, operator training, and non-normal procedure design. However, design is typically the biggest driver for shaping reliable operator performance over the long term.

It is also important to note that it is difficult to extract data on how frequently these alerts are handled well, and, thus, we cannot know the overall level of flight crew reliability. Certainly, the fatal accident rate is quite low, but while significant failures are relatively rare, there is value in a thorough analysis of failures in alerting system design. For example, it is noteworthy that for 17 of the 46 events analyzed, there was no alert for the hazard that led to the accident. For the 25 events for which there was hazard alerting, the reasons for failure were varied, covering much of the full range of the performance cycle (Rasmussen, 1983).

The evolution in alerting across generations of airplanes is an important element of the story.

The Evolution of Alerts

An airplane can be in service 30 years or more and older technology remains in the operating fleet even when new innovations improve safety. Alerting schemes, like many other elements of the flight deck interface, have evolved over time and alerts were sometimes added in a piecemeal manner. In fact, many of the airplanes involved in the events where there was no alerting were airplanes designed pre-1990.

For Boeing airplanes, there was a significant advance in alerting in the early 1980s (with the 757/767 joint entry into service) with the introduction of a centralized alerting system (the EICAS display in Boeing airplanes; the ECAM display in Airbus airplanes). Centralized alerting meant that, instead of alerts being distributed around the flight deck interface, they were presented on a single display placed in a forward area visible to both pilots. This change was in response to a proliferation of airplane system alerts in the previous generation (see Boucek et al., 1977).

However, this central alerting advance was applied only to airplane system failures, such as hydraulic system or electrical system failures. Flight path management hazards, on the other hand, continued to evolve without the level of integration brought to the airplane system alerts. Flight path-related alerts such as TCAS RAs, windshear, ground proximity, low airspeed, bank angle, and stall warning are related in that they all require the pilot to make flight control inputs. However, in even the newest generation of jet transports, pilots are alerted to these conditions using a mix of voices, sounds, master lights, visual alerts, and visual situation information (typically communicated through the Primary Flight Display (PFD), Navigational Display (ND) or other, more isolated lights near those displays). There are also differences in how guidance information is provided, or even if flight control guidance is provided.

For Boeing airplanes, the following describes the time course of alert evolution for a sample of alerts:

- Cabin altitude, which should be alerted at the warning level (two modalities), was alerted as a horn only in the 737 models and the early 747s. For the 737, it was 2011 before a light was added to help distinguish this alert from the Take-off Configuration alert (motivated by a 2005 accident).
- There was initially no alerting for bank angle (exceeding 35° of roll) in airplanes that went into service prior to 1989. Most of these airplanes were offered an option to install an aural only alert (for a caution-level alert) in 1987. After 1989, airplanes entering service had the necessary aural (voice) and visual (PFD indications). An alert with flight control guidance was not developed until 2016 (see Ewbank et al., 2016).
- There was initially no alerting for low airspeed, which should be a caution-level alert, on the early 737s. For airplanes that went into service in 1982 (but also the later 737 NG in 1998), there was only a

visual indication on the airspeed indicator on the PFD. The appropriate two-modality scheme arrived in 1989 with the 747-400.
- The approach to stall alert has always been a two-modality alert, from the very first Boeing 737 models. The early airplanes (737 and 747 classics) presented just the stick shaker (tactile and aural, and salient). PFD indications (visual) were added in 1982.
- Unreliable airspeed has always been a difficult situation to alert. There was no alert at all until 1989 for this caution-level alert. Some visual alerting was provided in the mid-1990s. The 737 Classic and NG never went beyond this visual-only scheme. True two-modality alerting arrived in 1997 for the other airplanes.

The lesson to be drawn from these time histories of individual alerts is that flight path management alerting has evolved in a fragmented manner; it has yet to be integrated like airplane system failures were. A significant contributor to this outcome was a belief that pilots were monitoring flight path effectively; that is, pilots were believed to be on top of airspeed, bank angle, angle of attack, and unreliable airspeed. It was assumed that pilots would notice these deviations and manage the situation appropriately. Accidents have shown this assumption to be generous for many operations, and recently, there is a new awareness regarding the weaknesses in pilot monitoring (e.g., CAA, 2013). Alerting schemes for flight path management have not yet evolved to reflect this new understanding.

Conclusions

This review of a large, representative set of recent commercial aviation accidents revealed that, for the operational fleet, there are still situations in which the alerting system fails to lead to successful management of the hazard. A failure in alerting (broadly defined here) contributed to 42 of 46 accidents!
Some of the alerting issues uncovered by this review have been addressed in newer airplanes through improved technology (e.g., unreliable airspeed). However, airplanes with older technology will remain in the operating fleet for years to come and vulnerabilities remain.
Only be studying the performance of alerting systems in the events in which they failed can we identify the types of alerting system design changes that are needed. A significant factor to consider in this evaluation is that flight path management alerting is a product of a less integrated approach to alerting than is airplane system alerting. One reason for this more

piecemeal approach was the set of assumptions about the effectiveness of pilot monitoring of flight path.

References

Boucek, G.P., Veitengruber, J.E., & Smith, W.D. (1977). Aircraft alerting systems criteria study. Vol II: Human factors guidelines for aircraft alerting systems. FAA Final Report FAA-RD-76-222. Washington, DC: FAA.

Civil Aviation Authority (2013). Monitoring matters. CAA Paper 2013/02.

Commercial Aviation Safety Team (2014). Airplane State Awareness JSAT Final Report

Dehais, F., Causse, M. Vachon, F., Regis, N., Menant, E., & Tremblay, S. (2014). Failure to detect critical auditory alerts in the cockpit: Evidence for inattentional

Ewbank, C.E., Mumaw, R.J., & Snow, M.P. (2016). Development of the enhanced bank angle warning. Proceedings of the IEEE/AIAA 35th Digital Avionics Systems Conference (DASC), Sacramento, CA.

Mumaw, R. (2017). Analysis of alerting system failures in commercial aviation accidents. NASA Technical report.

Rasmussen, J. (1983). Skills, rules, and knowledge; signals, signs, and symbols and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13 (3), pp. 257-266.